Generating Personalized Dialogue via Multi-Task Meta-Learning

Lee Jing Yang[1], Lee Kong Aik[2], Gan Woon Seng[3]

School of Electrical and Electronic Engineering, Nanyang Technological University, Singapore[1,3]

Institute for Infocomm Research (I²R), A*STAR, Singapore[2]

20th September 2021

Agenda

- Personalized Dialogue Generation
  - PersonaChat Corpus
- Meta-Learning
  - Model Agnostic Meta-Learning (MAML)
- Multi-task Meta-Learning
  - Multi-Task Meta-Learning (MTML)
  - Alternating Multi-Task Meta-Learning (AMTML)
- Implementation
- Results & Discussion
- Conclusion

Personalized Dialogue Generation

- Personalized dialogue generation involves the generation of dialogue responses which incorporate the personality of the dialogue agent.
- Typically, generated response is conditioned on both the persona information and dialogue context (dialogue history).
  - Persona information could constitute textual persona descriptions or profile attributes
- In a real-world scenario, persona information is not readily available.
- Condition the generated response on past dialogue examples instead.
  - However, this requires a large amount of persona-specific dialogue examples.
- Task: Generate personalized dialogue with small amount of persona-specific dialogue examples i.e., *few-shot personalized dialogue generation.*
- Hence, we propose a multi-task meta-learning framework to address the task of few-shot personalized dialogue generation.

PersonaChat Corpus

- In the PersonChat Corpus (Zhang et al., 2018), each conversation consists of two interlocutors trying to find out more about each other.
- The PersonaChat corpus comprises 1,155 distinct personas, each consisting of several persona statements.
- The dialogue agent must generate fluent, coherent dialogue consistent with the corresponding persona description, which is defined by several persona statements.

Persona Statements

I read twenty books a year.
I'm a stunt double in my second job.
I only eat kosher.
I was raised in a single parent household.

Response Label

I have a lot of books to read

Dialogue Context

Human: Hello what are you doing today?
Agent: i am good , i just got off work and tired , i have two jobs.
Human: i just got done watching a horror movie.
Agent: i rather read , i've read about 20 books this year
Human: wow! i do love a good horror movie. loving this cooler weather

Meta-Learning

- **Goal**: Train a model capable of rapid and efficient generalization to unseen tasks (few-shot learning).
  - Learning to learn
- Generally, meta-learning models are trained on many tasks and are then evaluated on their ability to adapt to new, unseen tasks.
- Types of meta-learning:
  - **Metric-based**
    - Involves learning a metric or distance function over tasks.
    - Eg. Convolutional Siamese Neural Network (Koch et al., 2015), Memory Augmented Neural Network (Santoro et al., 2016)
  - **Model-based**
    - Involves a specific model architecture.
    - Eg. Prototypical Networks (Snell et al., 2016), Matching Networks (Vinyals et al., 2017), Simple Neural Attentive Meta Learner (Mishra et al., 2018)
  - **Optimization-based**
    - Involves modifying the optimization algorithm.
    - Eg. Model Agnostic Meta Learning (Finn et al., 2017), Reptile (Nichol et al., 2018)

Model-Agnostic Meta-Learning (MAML)

- MAML aims to learn a parameter initialization capable of rapid and efficient adaptation.
- Most popular meta-learning framework in recent years.
- For each iteration i:
  - Sample tasks from the training set $\mathcal{D}_{train}$ to form a query set $\mathcal{O}^i$ and a support set $\mathcal{T}^i$.
  - Update parameters using the support set (inner loop):

$$\phi' = \phi - \eta_t \nabla_\phi L^{\mathcal{T}^i}(\phi)$$

  - Parameters are trained by optimizing for model performance with respect to ϕ across tasks in query set. This results in the following objective function:

$$\min_\phi \sum_{\mathcal{C}^i \sim \mathcal{D}_{train}} L^{\mathcal{O}^i}(\phi') = \sum_{\mathcal{C}^i \sim \mathcal{D}_{train}} L^{\mathcal{O}^i}(\phi - \eta_t \nabla_\phi L^{\mathcal{T}^i}(\phi))$$

  - Update parameters using the query set (outer loop):

$$\phi = \phi - \eta_o \nabla_\phi \sum_{\mathcal{O}^i} L^{\mathcal{O}^i}(\phi')$$

- Persona Agnostic Meta Learning (PAML) (Madotto et al, 2019)
  - Application of MAML to personalized dialogue generation.
  - Each unique persona is regarded as a unique task, and model is trained on a distribution of personas.

Multi-Task Meta-Learning

- We introduce a persona reconstruction task, which would result in parameters which induces the persona from the dialogue context to a larger extent.
- Persona reconstruction will be introduced via a persona reconstruction loss, which involves generating all corresponding persona statements in its entirety given the dialogue context.
- Given parameters ϕ, dialogue context $x_{1:t-1}$ and persona statements $\mathcal{P}_{1:N}$, the loss functions are defined as follows:

Response Generation Loss

$$f_\phi(x_t|x_{1:t-1}) = p(x_t|x_{1:t-1};\phi)$$

$$L_{res}(\phi) = -\sum_{t=1}^{T} log\, p(x_t|x_{1:t-1};\phi)$$

Persona Reconstruction Loss

$$\overline{\mathcal{P}} = concat(\mathcal{P}_{1:N})$$

$$f_\phi(\overline{\mathcal{P}}|x_{1:t-1}) = p(\overline{\mathcal{P}}|x_{1:t-1};\phi)$$

$$L_{rec}(\phi) = -\sum_{t=1}^{T} log\, p(\overline{\mathcal{P}}|x_{1:t-1};\phi)$$

- Persona reconstruction is introduced *only during meta-learning.*
- We introduce two multi-task meta-learning frameworks:
  - Multi-Task Meta-Learning (MTML)
  - Alternating Multi-Task Meta-Learning (AMTML)

Multi-Task Meta-Learning (MTML) Framework

- Combine the response generation loss and the persona reconstruction loss to form multi-task loss function.
- Hyperparamter α determines the contribution of the persona reconstruction and response generation loss.
- For each iteration i:
  - Update parameters using the support set

$$L^{\mathcal{T}^i}(\phi) = \alpha L_{res}^{\mathcal{T}^i}(\phi) + (1-\alpha) L_{rec}^{\mathcal{T}^i}(\phi)$$

$$\phi' = \phi - \eta_t \nabla_\phi L^{\mathcal{T}^i}(\phi)$$

  - Update parameters using the query set

$$L^{\mathcal{O}^i}(\phi') = \alpha L_{res}^{\mathcal{O}^i}(\phi') + (1-\alpha) L_{rec}^{\mathcal{O}^i}(\phi')$$

$$\phi = \phi - \eta_o \nabla_\phi \frac{1}{M} \sum_{\mathcal{O}^i} L^{\mathcal{O}^i}(\phi')$$

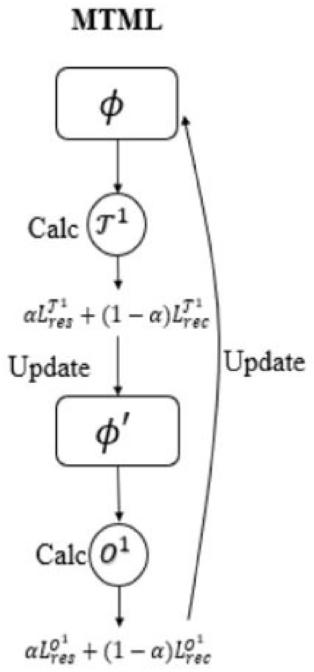

Fig 1: Overview of MTML

- During inference, dialogue examples are used to finetune the model via only the response generation loss. Persona statements are not required during inference.

Alternating Multi-Task Meta-Learning (AMTML) Framework

- Constantly alternate between response generation loss L_{res} and persona reconstruction loss L_{rec}.

- At every iteration, the inner loop and outer loop utilize different loss functions.

- For each iteration *i*:
  - Update parameters using the support set with either the response generation loss and persona reconstruction loss

    $\phi' = \phi - \eta_t \nabla_\phi L^{\mathcal{T}^i}_{rec}(\phi)$ or $\phi' = \phi - \eta_t \nabla_\phi L^{\mathcal{T}^i}_{res}(\phi)$

  - Update parameters using the query set with the alternate loss function

    $L^{\mathcal{O}^i}(\phi') = L^{\mathcal{O}^i}_{res}(\phi')$ or $L^{\mathcal{O}^i}(\phi') = L^{\mathcal{O}^i}_{rec}(\phi')$

    $$\phi = \phi - \eta_o \nabla_\phi \frac{1}{M} \sum_{\mathcal{O}^i} L^{\mathcal{O}^i}(\phi')$$

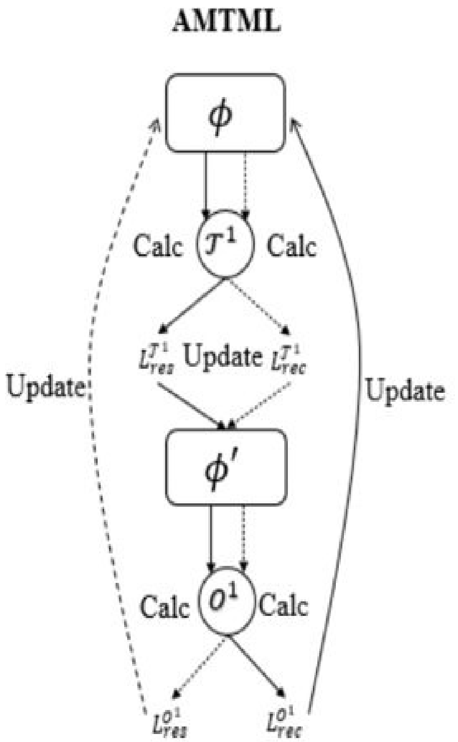

Fig 2: Overview of AMTML

- During inference, dialogue examples are used to finetune the model via only the response generation loss. Persona statements are not required during inference.

Experiment

- Corpus: PersonaChat

- Base Model: Transformer (6 encoder layers, 6 decoder layers and 4 attention heads) with Glove embedding

- Evaluation:
  - Human Evaluation: Consistency, Fluency, Coherence
  - Automatic Metrics: PPL, BLEU, C-score (Madotto et al., 2019)
    - C-score is derived via a BERT-based Natural Language Inference (NLI) model finetuned to indicate if the response entails, contradicts, or is neutral with respect to the corresponding persona statements.
    - For a given response, a score of 1, -1 and 0 is assigned if the response entails, contradicts or is neutral with regard to each of the persona statements in the persona description.
    - The final C-score for each response is attained by summing the scores for all persona statements in the description.

- Implemented models:
  - ***Std***: standard Transformer model with only the dialogue context as input
  - ***Std_p***: standard Transformer model with both the dialogue context and persona statements as inputs.
  - ***PAML***: standard Transformer pretrained via PAML (Madotto et al., 2019)
  - ***$MTML_\alpha$*** : standard Transformer pretrained via MTML where α = [0.9, 0.8, 0.7, 0.6, 0.5]
  - ***AMTML***: a standard Transformer pretrained via AMTML

- Experiments were conducted in both the 10-shot and 5-shot setting.

Results & Discussion

Table 1: 10-shot automatic evaluation results

	PPL	BLEU	C-score
Std	35.87	0.93	0.00
Std_p	38.72	1.66	0.10
$PAML$	41.80	0.71	0.19
$MTML_{0.5}$	77.32	0.53	0.46
$MTML_{0.6}$	57.10	0.53	0.41
$MTML_{0.7}$	52.44	0.57	0.47
$MTML_{0.8}$	43.28	0.42	0.34
$MTML_{0.9}$	40.39	0.71	0.21
$AMTML$	48.66	0.48	0.29

Table 2: 10-shot automatic evaluation results

	Consistency	Fluency	Coherence
Std	0.10	0.87	0.20
Std_p	0.09	0.89	0.21
$PAML$	0.13	0.84	0.27
$MTML_{0.5}$	0.22	0.03	-0.12
$MTML_{0.6}$	0.24	0.65	0.15
$MTML_{0.7}$	0.29	0.65	0.13
$MTML_{0.8}$	0.22	0.78	0.29
$MTML_{0.9}$	0.15	0.77	0.19
$AMTML$	0.23	0.85	0.20

- For ***MTML***$_{\alpha}$, optimal value of α = 0.8
  - Decreasing α would generally increase the Consistency/C-score at the expense of fluency and coherence.
- ***MTML***$_{0.8}$ and ***AMTML*** increases the amount of persona information reflected in the generated dialogue without compromising the coherence and fluency.
- In general, there is still room for improvement regarding contextual coherence.
- Results in the five-shot setting generally lower but follow the same trend. Full results are provided in the paper.

Results and Discussion

Table 3: Human and automatic evaluation results of ***P²Bot***

Automatic	PPL	BLEU	C-score
P^2Bot	18.1	0.61	0.33
Human	Consistency	Fluency	Coherence
P^2Bot	0.39	0.91	0.43

- ***P²Bot*** involves finetuning the GPT pretrained language model on the training set via a transmitter receiver framework.
- For ***P²Bot***, persona statements are provided to the model along with the dialogue context during inference.
- In a 10-shot setting, ***MTML$_{0.8}$*** and ***AMTML*** achieved comparable C-score/Consistency despite *not* being presented with the persona statements during inference.
- However, ***P²Bot*** still has the edge in terms of general fluency and coherence.

Conclusion

- For the task of few-shot personalized dialogue generation, both MTML and AMTML effectively increases the amount of persona information reflected in the generated dialogue responses compared to prior work.
- The overall fluency and contextual coherence of the generated responses can still be improved.
- Future work could involve incorporating pretrained language models (T5, GPT-2, BERT...) in a meta-learning framework.
  - However, utilizing the standard MAML framework might be challenging due to the sheer size of pretrained language models.
  - First order meta-learning frameworks such as First Order MAML (FOMAML) or Reptile can be considered instead.

References

- Jake Snell, Kevin Swersky, and Richard Zemel. 2017. Prototypical networks for few-shot learning. In Advances in Neural Information Processing Systems, volume 30. Curran Associates, Inc.
- Oriol Vinyals, Charles Blundell, Timothy Lillicrap, koray kavukcuoglu, and Daan Wierstra. 2016. Matching networks for one shot learning. In Advances in Neural Information Processing Systems, volume 29.Curran Associates, Inc.
- Adam Santoro, Sergey Bartunov, Matthew Botvinick, Daan Wierstra, and Timothy Lillicrap. 2016. Meta-learning with memory-augmented neural networks. In Proceedings of The 33rd International Conference on Machine Learning, volume 48 of Proceedings of Machine Learning Research, pages 1842–1850, New York, New York, USA. PMLR.
- N. Mishra, Mostafa Rohaninejad, Xi Chen, and P. Abbeel. 2018. A simple neural attentive meta-learner. In ICLR.
- Chelsea Finn, Pieter Abbeel, and Sergey Levine. 2017.Model-agnostic meta-learning for fast adaptation of deep networks. In Proceedings of the 34th International Conference on Machine Learning, volume 70 of Proceedings of Machine Learning Research, pages 1126–1135, International Convention Centre, Sydney, Australia. PMLR.
- Alex Nichol, Joshua Achiam, and John Schulman.2018.On first-order meta-learning algorithms. ArXiv, abs/1803.02999
- Andrea Madotto, Zhaojiang Lin, Chien-Sheng Wu, and Pascale Fung. 2019. Personalizing dialogue agents via meta-learning. In Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics, pages 5454–5459, Florence, Italy. Association for Computational Linguistics.
- Saizheng Zhang, Emily Dinan, Jack Urbanek, ArthurSzlam, Douwe Kiela, and Jason Weston. 2018. Personalizing dialogue agents: I have a dog, do you have pets too? In Proceedings of the 56th An-nual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), pages 2204–2213, Melbourne, Australia. Association for Computational Linguistics.

Thank You